

On Bootstrap Averaging Empirical Bayes Estimators*

SHONOSUKE SUGASAWA

Risk Analysis Research Center, The Institute of Statistical Mathematics

Abstract. Parametric empirical Bayes (EB) estimators have been widely used in variety of fields including small area estimation, disease mapping. Since EB estimator is constructed by plugging in the estimator of parameters in prior distributions, it might perform poorly if the estimator of parameters is unstable. This can happen when the number of samples are small or moderate. This paper suggests bootstrapping averaging approach, known as “bagging” in machine learning literatures, to improve the performances of EB estimators. We consider two typical hierarchical models, two-stage normal hierarchical model and Poisson-gamma model, and compare the proposed method with the classical parametric EB method through simulation and empirical studies.

Key words: Bagging; Hierarchical model; Mean squared error; Poisson-gamma model

1 Introduction

The parametric empirical Bayes estimators (Morris, 1983) are known to be a useful method producing reliable estimates of multidimensional parameters. This technique is widely used in variety of fields such as small area estimation (Rao and Molina, 2015) and disease mapping (Lawson, 2013). Let $\theta_1, \dots, \theta_m$ be the multiple parameters of interest, and y_1, \dots, y_m be the independent observations generated from the distribution $f_i(y_i|\theta_i)$, $i = 1, \dots, m$. To carry out an empirical Bayes estimation, it is assumed that the parameters $\theta_1, \dots, \theta_m$ independently follows the distribution $g(\theta_i; \phi)$, where ϕ is a vector of unknown parameters. Therefore, we obtain a two stage model:

$$y_i|\theta_i \sim f_i(y_i|\theta_i), \quad \theta_i \sim g(\theta_i; \phi), \quad i = 1, \dots, m, \quad (1)$$

which are independent for $i = 1, \dots, m$. Under the setting, the posterior distribution of θ_i is given by

$$\pi(\theta_i|y_i; \phi) = \frac{f_i(y_i|\theta_i)g(\theta_i; \phi)}{\int f_i(y_i|\theta_i)g(\theta_i; \phi)d\theta_i}, \quad i = 1, \dots, m.$$

The Bayes estimator $\tilde{\theta}_i$ of θ_i under squared error loss is the conditional expectation (posterior mean) of θ_i given y_i , that is

$$\tilde{\theta}_i \equiv E[\theta_i|y_i; \phi] = \frac{\int \theta_i f_i(y_i|\theta_i)g(\theta_i; \phi)d\theta_i}{\int f_i(y_i|\theta_i)g(\theta_i; \phi)d\theta_i}, \quad i = 1, \dots, m. \quad (2)$$

*This version: April 28, 2017

However, the Bayes estimator $\tilde{\theta}_i$ depends on unknown model parameters ϕ , which can be estimated from the marginal distribution of all the data $y = \{y_1, \dots, y_m\}$, given by

$$L(\phi) = \prod_{i=1}^m \int f_i(y_i|\theta_i)g(\theta_i; \phi)d\theta_i.$$

Using the marginal distribution of y , one can immediately define the maximum likelihood (ML) estimator as the maximizer of $L(\phi)$. Based on the estimator $\hat{\phi}$, we obtain the empirical Bayes (EB) estimator of θ_i as $\hat{\theta}_i = E[\theta_i|y_i; \hat{\phi}]$.

The variability of the EB estimator $\hat{\theta}_i$ can be measured by the integrated mean squared error (MSE) $E[(\hat{\theta}_i - \theta_i)^2]$, where the expectation is taken with respect to θ_i 's and y_i 's following the model (1). Since $\tilde{\theta}_i$ is the conditional expectation as given in (2), the MSE can be decomposed as $E[(\hat{\theta}_i - \theta_i)^2] = R_1 + R_2$ with $R_1 = E[(\tilde{\theta}_i - \theta_i)^2]$ and $R_2 = E[(\hat{\theta}_i - \tilde{\theta}_i)^2]$. The first term R_1 is not affected by the estimation of ϕ whereas the second term R_2 reflects the variability of the ML estimator $\hat{\phi}$, so that the second term can be negligibly small when m is large. However, in many applications, m might be small or moderate, in which the contribution of the second term to the MSE cannot be ignored. Hence, the EB estimator might perform poorly depending on the ML estimator $\hat{\phi}$. To overcome this problem, we propose to use the bootstrap averaging technique, known as ‘‘bagging’’ (Breiman, 1996) in machine learning literatures. This method produces many estimators based on bootstrap samples, and average them to produce a stable estimator. We adapt the bagging method to the EB estimation to improve the performances of EB estimators under small or moderate m .

This paper is organized as follows: In Section 2, we consider mean squared errors of EB estimators and propose a bootstrap averaging empirical Bayes (BEB) estimator for decreasing the mean squared error. In Section 3 and Section 4, we apply the BEB estimators in well-known two-stage normal hierarchical model and Poisson-gamma model, respectively, and compare the performances between BEB and EB estimators through simulation and empirical studies. In Section 5, we provide conclusions and discussions.

2 Bootstrap Averaging Empirical Bayes Estimators

As noted in the previous section, the performances of the EB estimators depend on the variability of the estimator $\hat{\phi}$, which cannot be ignored when m is not large. To reduce the variability of the empirical Bayes estimator $\hat{\theta}_i$, we propose to average many empirical Bayes estimators with bootstrap estimates of ϕ rather than computing one empirical Bayes estimator from the observation $Y = \{y_1, \dots, y_m\}$. Specifically, letting $Y_{(b)} = \{y_1^{(b)}, \dots, y_m^{(b)}\}$ be a bootstrap samples of the original observation Y , we define $\hat{\phi}_{(b)}$ be an estimator of ϕ based on the bootstrap sample $Y_{(b)}$. Then the bagging empirical Bayes (BEB) estimator is given by

$$\hat{\theta}_i^{\text{Boot}} = \frac{1}{B} \sum_{b=1}^B \tilde{\theta}_i(y_i, \hat{\phi}_{(b)}). \quad (3)$$

Similarly to Breiman (1996), we note that

$$\begin{aligned} \frac{1}{B} \sum_{b=1}^B \left\{ \tilde{\theta}_i(y_i, \hat{\phi}_{(b)}) - \theta_i \right\}^2 &= \frac{1}{B} \sum_{b=1}^B \tilde{\theta}_i(y_i, \hat{\phi}_{(b)})^2 - 2\hat{\theta}_i^{\text{Boot}} \theta_i + \theta_i^2 \\ &\geq \left\{ \frac{1}{B} \sum_{b=1}^B \tilde{\theta}_i(y_i, \hat{\phi}_{(b)}) \right\}^2 - 2\hat{\theta}_i^{\text{Boot}} \theta_i + \theta_i^2 = (\hat{\theta}_i^{\text{Boot}} - \theta_i)^2. \end{aligned}$$

By taking expectation with respect to the model (1), we have

$$\frac{1}{B} \sum_{b=1}^B \mathbb{E} \left[\left\{ \tilde{\theta}_i(y_i, \hat{\phi}_{(b)}) - \theta_i \right\}^2 \right] \geq \mathbb{E} \left[(\hat{\theta}_i^{\text{Boot}} - \theta_i)^2 \right],$$

which means that the integrated MSE of BEB estimator (3) is smaller than bootstrap average of the integrated MSE of the EB estimator. Hence, the BEB estimator is expected to perform better than the EB estimator. The amount of improvement depends on

$$\frac{1}{B} \sum_{b=1}^B \tilde{\theta}_i(y_i, \hat{\phi}_{(b)})^2 - \left\{ \frac{1}{B} \sum_{b=1}^B \tilde{\theta}_i(y_i, \hat{\phi}_{(b)}) \right\}^2 = \frac{1}{B} \sum_{b=1}^B \left\{ \tilde{\theta}_i(y_i, \hat{\phi}_{(b)}) - \hat{\theta}_i^{\text{Boot}} \right\}^2,$$

which is the bootstrap variance of the EB estimator and it vanishes as $m \rightarrow \infty$ but it would not be negligible when m is not large. Therefore, when m is small or moderate, the BEB estimator would improve the performance of the EB estimator. In the subsequent section, we investigate the performances of the EBE estimator compared with the EB estimator in the widely-used hierarchical models.

3 Two-stage normal hierarchical model

3.1 Model description

We first consider the two-stage normal hierarchical model to demonstrate the proposed bagging procedure. The two-stage normal hierarchical model is described as

$$y_i | \theta_i \sim N(\theta_i, D_i), \quad \theta_i \sim N(\mathbf{x}_i^t \boldsymbol{\beta}, A), \quad i = 1, \dots, m, \quad (4)$$

where D_i is known sampling variance, \mathbf{x}_i and $\boldsymbol{\beta}$ are a vector of covariates and regression coefficients, respectively, A is an unknown variance. Let $\boldsymbol{\phi} = (\boldsymbol{\beta}^t, A)^t$ be the vector of unknown parameters. The model (6) is known as the Fay-Herriot model (Fay and Herriot, 1979) in the context of small area estimation.

Under the model (6), the Bayes estimator of θ_i is

$$\tilde{\theta}_i(y_i; \boldsymbol{\phi}) = \mathbf{x}_i^t \boldsymbol{\beta} + \frac{D_i}{A + D_i} (y_i - \mathbf{x}_i^t \boldsymbol{\beta}).$$

Concerning the estimation of unknown parameter $\boldsymbol{\phi}$, we here consider the maximum likelihood estimator for simplicity. Since $y_i \sim N(\mathbf{x}_i^t \boldsymbol{\beta}, A + D_i)$ under the model (6), the maximum likelihood estimator $\hat{\boldsymbol{\phi}}$ is defined as the maximizer of the function:

$$Q(\boldsymbol{\phi}) = \sum_{i=1}^m \log(A + D_i) + \sum_{i=1}^m \frac{(y_i - \mathbf{x}_i^t \boldsymbol{\beta})^2}{A + D_i}.$$

While several other estimating methods are available, we here only consider the maximum likelihood estimator for presentational simplicity. Using the maximum likelihood estimator $\hat{\phi}$, we obtain the EB estimator of θ_i as $\tilde{\theta}_i(y_i; \hat{\phi})$.

3.2 Simulation study

We here evaluate the performances of the BEB estimator together with the EB estimator under the normal hierarchical model (6) without covariates, namely $\mathbf{x}_i^t \boldsymbol{\beta} = \mu$. We considered $m = 10, 15, \dots, 40$. For each m , we set D_i as equally spaced points between 0.5 and 1.5. Concerning the true parameter values, we used $\mu = 0$ and four cases for A , namely $A = 0.1, 0.3, 0.5$ and 0.7 . The simulated data was generated from the model (6) in each iteration, and computed the EB and BEB estimates of θ_i . Based on $R = 5000$ simulation runs we calculated the simulated mean squared errors (MSE) defined as

$$\text{MSE} = \frac{1}{mR} \sum_{i=1}^m \sum_{r=1}^R (\hat{\theta}_i^{(r)} - \theta_i^{(r)})^2, \quad (5)$$

where $\hat{\theta}_i^{(r)}$ is the EBE or EB estimates and $\theta_i^{(r)}$ is the true value of θ_i in the r th iteration.

In Figure 1, we present the simulated MSE of the EB estimator as well as the three BEB estimator using 25, 50 and 100 bootstrap samples under various settings of A and m . It is observed that the BEB estimator performs better than the EB estimator on the whole. In particular, the improvement is greater when A is small compared with D_i , which is often arisen in practice. Moreover, as the number of m gets larger, the MSE differences get smaller since the variability of estimating ϕ vanishes when m is sufficiently large. We also found that the ML estimator of A often produces 0 estimates when m is small, in which the EB estimator is known to perform poorly. However, the BEB estimator can avoid the problem since the BEB estimator is aggregated by B bootstrap estimators and at least one bootstrap estimates should be non-zero. In fact, by investigating the case where the ML estimator produces 0 estimates of A , the some bootstrap estimates of A were away from 0. This would be one of the reason why the BEB estimator performs better than the EB estimator in this setting.

3.3 Example: corn data

We next illustrate the performances of the BEB estimator by using the corn and soybean productions in 12 Iowa counties, which has been used as an example in the context of small area estimation. Especially, we use the area-level data set given in table 6 in Dass et al. (2012) and we here focus only on corn productions for simplicity. The data set consists of $m = 8$ areas with sample sizes in each area ranging from 3 to 5, and survey data of corn production y_i , sampling variance D_i and the satellite data of corn x_i as the covariate observed in each area. We considered the following hierarchical model:

$$y_i | \theta_i \sim N(\theta_i, D_i), \quad \theta_i \sim N(\beta_0 + \beta_1 x_i, A), \quad i = 1, \dots, m, \quad (6)$$

where β_0, β_1 and A are unknown parameters. For the data set, we computed the BEB as well as EB estimators. We used 1000 bootstrap samples for computing the

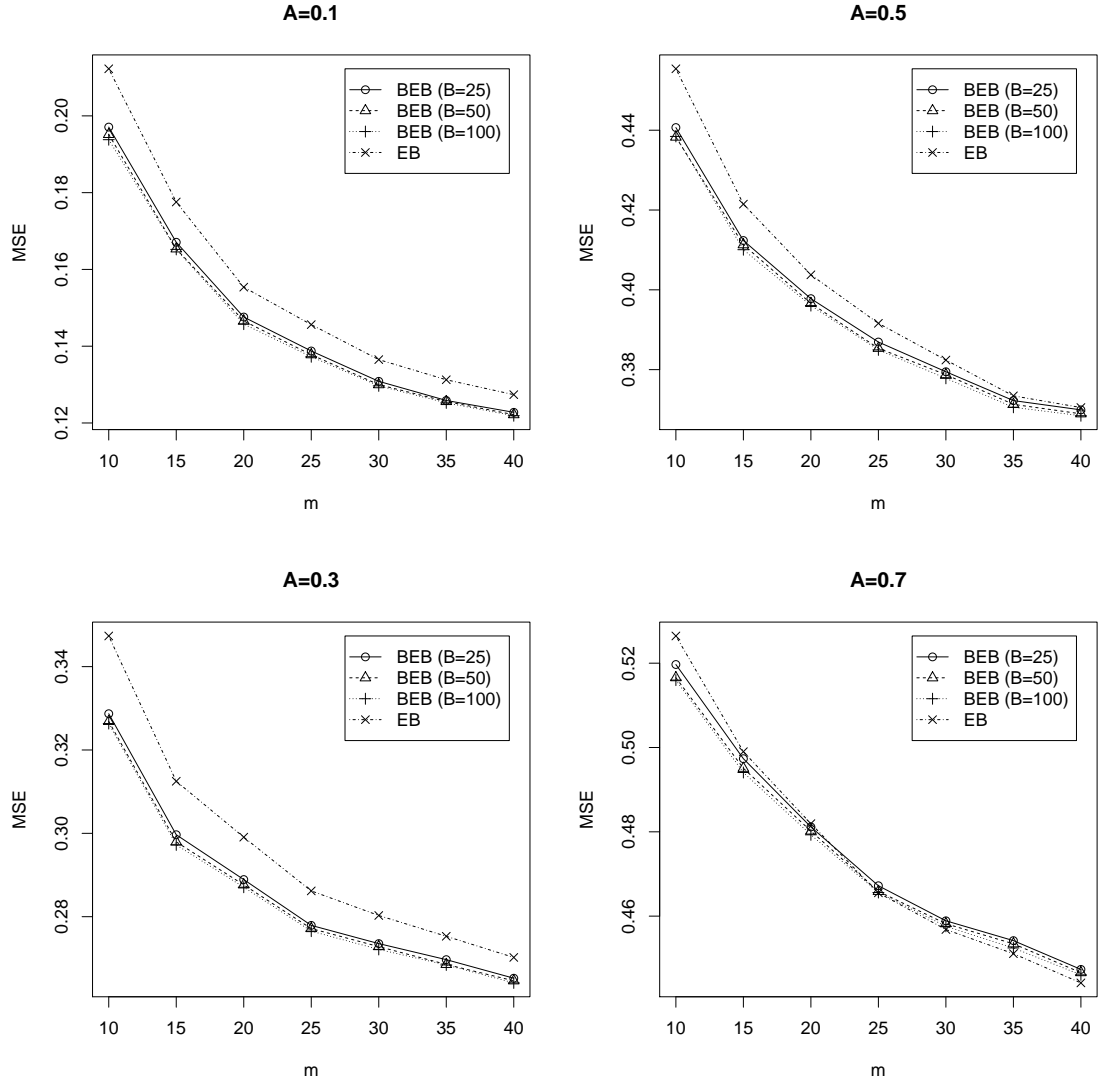


Figure 1: The simulated MSE of three estimators, BEB (bootstrap averaging empirical Bayes estimator) and EB (empirical Bayes estimator) in two-stage normal hierarchical model.

BEB estimator. In Figure 2, we present the histogram of the bootstrap estimates used in the BEB estimates and the maximum likelihood (ML) estimates used in the EB estimators. We can observe that the bootstrap estimates vary depending on the bootstrap samples. Moreover, in Table 1, we show the BEB and EB estimates of θ_i , which shows that the BEB estimator produces different estimates from the EB estimator since the number of areas m is only 8.

Table 1: Direct estimates (DE) y_i , standard deviation (SV) $\sqrt{D_i}$, empirical Bayes (EB) estimates and bagging empirical Bayes (BEB) estimates in each county.

	DE	SD	EB	BEB
Franklin	158.62	5.70	155.79	141.08
Pocahontas	102.52	43.41	102.82	97.48
Winnebago	112.77	30.55	119.74	117.34
Wright	144.30	54.00	127.86	124.05
Webster	117.59	21.30	109.61	102.42
Hancock	109.38	15.66	121.84	126.51
Kossuth	110.25	12.11	116.05	118.53
Hardin	120.05	36.81	136.97	137.05




Figure 2: The histograms of the bootstrap estimates of β_0 (left), β_1 (center) and A (right). Each vertical line denotes the maximum likelihood estimate.

4 Poisson-gamma model

4.1 Setup

The Poisson-gamma model (Clayton and Kalder, 1987) is described as

$$z_i | \theta_i \sim \text{Po}(n_i \theta_i), \quad \theta_i \sim \Gamma(\nu m_i, \nu), \quad i = 1, \dots, m, \quad (7)$$

where $m_i = \exp(\mathbf{x}_i^t \boldsymbol{\beta})$, \mathbf{x}_i and $\boldsymbol{\beta}$ are a vector of covariates and regression coefficients, respectively, ν is an unknown scale parameter. This model is used as the standard method of disease mapping. Let $\boldsymbol{\phi} = (\boldsymbol{\beta}^t, \nu)^t$ be the vector of unknown parameters. The model (9) is known as the Poisson-Gamma model considered in Clayton and Kaldor (1987) and used in disease mapping.

Under the model (9), the Bayes estimator of θ_i is given by

$$\tilde{\theta}_i(y_i; \boldsymbol{\phi}) = \frac{z_i + \nu m_i}{n_i + \nu}.$$

Since the Bayes estimator depends on unknown $\boldsymbol{\phi}$, we need to replace $\boldsymbol{\phi}$ by its estimator. Noting that the gamma prior of θ_i is a conjugate prior for the mean parameter in the Poisson distribution, the marginal distribution of y_i is the negative binomial distribution with the probability function:

$$f_m(y_i; \boldsymbol{\phi}) = \frac{\Gamma(z_i + \nu m_i)}{\Gamma(z_i + 1) \Gamma(\nu m_i)} \left(\frac{n_i}{n_i + \nu} \right)^{z_i} \left(\frac{\nu}{n_i + \nu} \right)^{\nu m_i}.$$

Then the maximum likelihood estimator of ϕ is defined as $\hat{\phi} = \operatorname{argmax}_{\phi} \sum_{i=1}^m \log f_m(y_i; \phi)$, which enables us to obtain the empirical Bayes estimator $\tilde{\theta}_i(y_i; \hat{\phi})$.

4.2 Simulation study

We next evaluated the performances of the BEB estimator under the Poisson-gamma model without covariates, described as

$$z_i | \theta_i \sim \text{Po}(n_i \theta_i), \quad \theta_i \sim \Gamma(\nu \mu, \nu), \quad i = 1, \dots, m, \quad (8)$$

where we set $\mu = 1$ and $\nu = 40, 60, 80$ and 100 . Note that ν is a scale parameter and $\text{Var}(\theta_i) = \mu/\nu$, so that random effect variance $\text{Var}(\theta_i)$ is a decreasing function of ν . Regarding the number of areas, we considered $m = 10, 15, \dots, 40$. For each m , we set n_i as rounded integers of equally spaced numbers between 10 and 50. Similarly to Section 3.2, using (5) with $R = 5000$ simulation runs, we calculated the MSE of the BEB estimator as well as the EB estimator of θ_i . The results are presented in Figure 3, which show that the BEB estimator tends to perform better than the EB estimator. In particular, the amount of improvement is greater when m is not large as we expected. Moreover, we can also observe that the MSE difference tends larger as ν gets larger, which corresponds to the case where the random effect variance gets smaller. This is consistent to the results in the normal model given in Section 3.2.

4.3 Example: Scottish lip cancer

We applied the BEB and EB method to the famous Scottish lip cancer data during the 6 years from 1975 to 1980 in each of the $m = 56$ counties of Scotland. For each county, the observed and expected number of cases are available, which are respectively denoted by z_i and n_i . Moreover, the proportion of the population employed in agriculture, fishing, or forestry is available for each county, thereby we used it as a covariate AFF_i , following Wakefield (2007). For each area, $i = 1, \dots, m$, we consider the Poisson-gamma model:

$$z_i | \theta_i \sim \text{Po}(n_i \theta_i), \quad \theta_i \sim \Gamma(\nu \exp(\beta_0 + \beta_1 \text{AFF}_i), \nu), \quad (9)$$

where θ_i is the true risk of lip cancer in the i th area, and β_0, β_1 and A are unknown parameters. For the data set, we computed the BEB as well as EB estimates of θ_i , where we used 1000 bootstrap samples for computing the BEB estimator. In Figure 4, we present the quantiles of the bootstrap estimates used in the BEB estimates and the maximum likelihood (ML) estimates used in the EB estimators. We can observe that the bootstrap estimates vary depending on the bootstrap samples while the variability seems small compared with Figure 2. This might come from that the number of areas in this case is much larger than the corn data in Section 3.3. Finally, in Figure 5, we show the scatter plot of percent relative difference between the BEB and EB estimates, that is, $100(\hat{\theta}_i^{\text{Boot}} - \hat{\theta}_i)/\hat{\theta}_i$, against the number of expected number of cases n_i . Figure 5 shows that the differences get larger as n_i gets small since the direct estimator $y_i = z_i/n_i$ of θ_i is shrunk toward the regression mean $\exp(\beta_0 + \beta_1 \text{AFF}_i)$ in areas with small n_i .

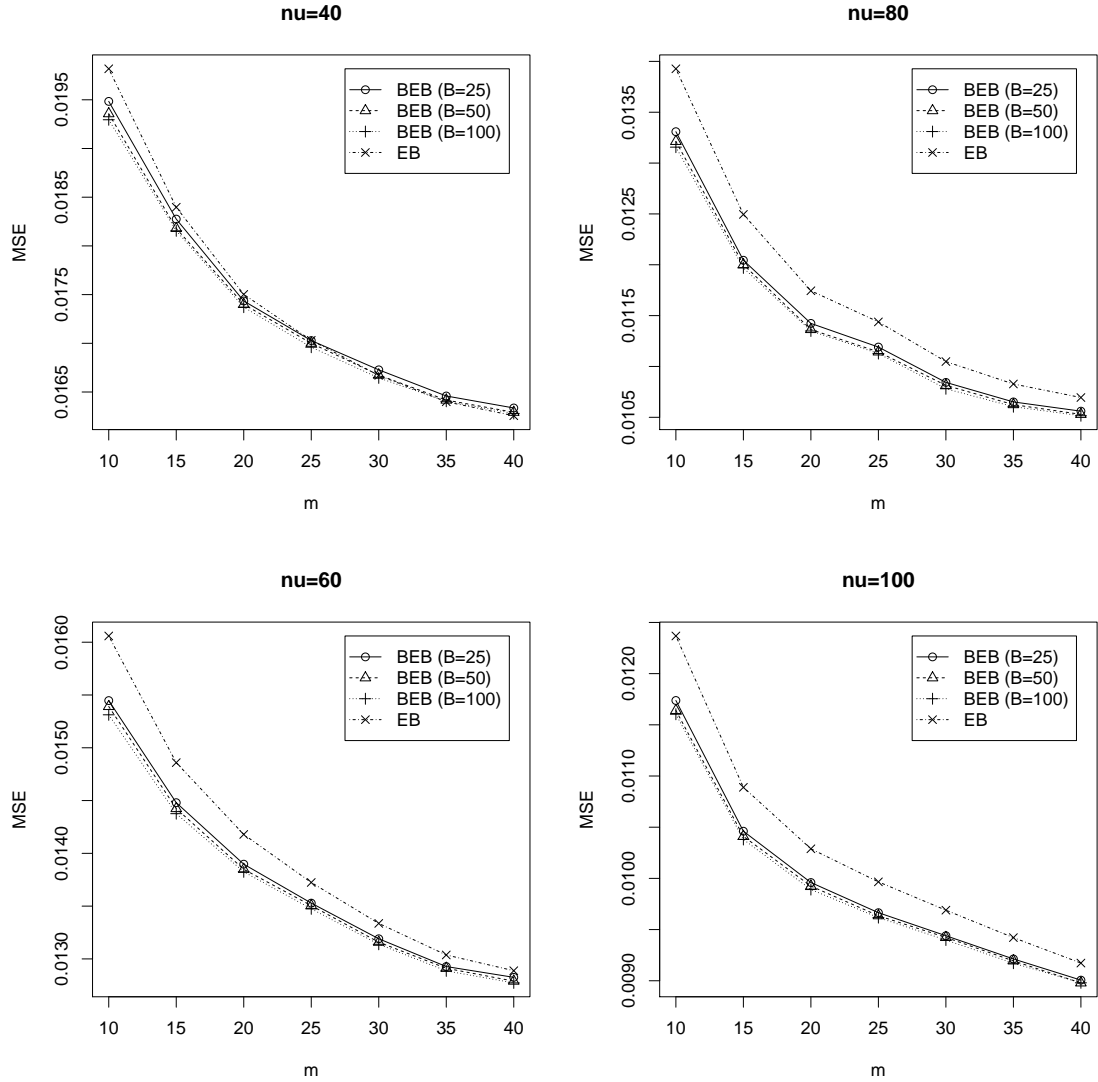


Figure 3: The simulated MSE of three estimators, BEB (bootstrap averaging empirical Bayes estimator) and EB (empirical Bayes estimator) in Poisson-gamma model.

5 Conclusion and Discussion

We have proposed the use of bootstrap averaging, known as “bagging” in the context of machine learning, for improving the performances of empirical Bayes (EB) estimators. We focused on two models extensively used in practice, two-stage normal hierarchical model and Poisson-gamma model. In both models, the simulation studies revealed that the bootstrap averaging EB (BEB) estimator performs better than the EB estimator.

In this paper, we considered the typical area-level models as an application of the BEB estimator. However, the BEB method would be extended to the more general case, for example, generalized linear mixed models. The detailed comparison in such models will be left to a future study.




Figure 4: The histograms of the bootstrap estimates of β_0 (left), β_1 (center) and ν (right). Each vertical line denotes the maximum likelihood estimate.




Figure 5: The scatter plot of the area sample size n_i against the percent relative difference between the BEB and EB estimates of θ_i .

References

- [1] Breiman, L. (1996). Bagging predictors. *Machine Learning*, **24**, 123-140.
- [2] Clayton, D. and Kaldor, J. (1987). Empirical Bayes estimates of age-standardized relative risks for use in disease mapping. *Biometrics*, **43**, 671-681.
- [3] Dass, S. C., Maiti, T., Ren, H. and Sinha, S. (2012). Confidence interval estimation of small area parameters shrinking both means and variances. *Survey Methodology*, **38**, 173-187.

- [4] Fay, R. and Herriot, R. (1979). Estimators of income for small area places: An application of James-Stein procedures to census. *Journal of the American Statistical Association*, **74**, 341-353.
- [5] Lawson, A. B. (2013). *Bayesian disease mapping: hierarchical modeling in spatial epidemiology, 2nd Edition*. Chapman and Hall/CRC press.
- [6] Morris, C. N. (1983). Parametric empirical Bayes inference: theory and applications. *Journal of the American Statistical Association*, **78**, 47-65.
- [7] Rao, J.N.K. and Molina, I. (2015) *Small Area Estimation, 2nd Edition*. Wiley.
- [8] Wakefield, J. (2007). Disease mapping and spatial regression with count data. *Biostatistics*, **8**, 158-183.